Optimality Theory and Human Sentence Processing:

Towards a Cross-Modular Analysis of Coordination

John C.J. Hoeks[1,2] & Petra Hendriks[1]

[1] Center for Language and Cognition Groningen

[2] BCN NeuroImaging Center

University of Groningen

Corresponding author:

John C.J. Hoeks

CLCG / BCN NeuroImaging Center

PO Box 716

9700 AS Groningen

The Netherlands

phone: +31(0)50 3637443

fax: +31(0)50 3636855

email: j.c.j.hoeks@rug.nl

Abstract

In this paper we propose a model of human sentence processing that is based on Optimality Theory (OT). In contrast to most other OT approaches to language processing, we use constraints from OT *semantics* rather than OT *syntax* to address on-line comprehension. We illustrate the workings of our model by investigating the processing of coordinated structures. The psycholinguistic evidence that is currently available suggests that the on-line comprehension of coordination is influenced by constraints from many different information sources: pragmatics, discourse semantics, lexical semantics, and syntax. The model we propose formalizes this cross-modular interaction of constraints, and yields concrete predictions with respect to both intermediate parsing preferences and final interpretations. Our ultimate aim is to develop a model of processing performance that at the same time is a fully functional model of linguistic competence.

KEYWORDS: Optimality Theory; grammar; transparent parser; sentence processing; on-line language comprehension; coordination; coordinated structures

1. Introduction

Optimality Theory (OT) is a powerful model of decision making in situations where there are multiple constraints pertaining to one or more alternative options. It was originally introduced as a model of linguistic *competence*, and as such has been successful in many linguistic domains: OT has been a standard theory in the field of phonology (e.g., Prince & Smolensky, 1993/2004), and it is influential in morphology (e.g., McCarthy & Prince, 1993), syntax (e.g., Barbosa, Fox, Hagstrom, McGinnis, & Pesetsky, 1998; Bresnan, 2000; Broekhuis & Vogel, 2009, 2010; Grimshaw, 1997; Legendre, Grimshaw, & Vikner, 2001; Sells, 2001; McCarthy, 2008), and semantics / pragmatics (e.g., Beaver, 2004; Blutner, 2000; de Hoop, Hendriks, & Blutner, 2007; de Hoop & de Swart, 2000; Hendriks & de Hoop, 2001; Hendriks & Spenader, 2004, 2005/2006; Zeevat, 2000). We will argue that, in addition to being an excellent framework for describing linguistic competence, OT is also very well suited as a model of linguistic *performance*, permitting us to model the interaction of many different information sources in an incremental fashion.

In most constraint-based models of language processing (e.g., models proposed by MacDonald, Pearlmutter, & Seidenberg, 1994, and by Trueswell & Tanenhaus, 1994), the interpretation of language input is conceived of as a process where many different, often probabilistic, factors provide support for one or the other syntactic structure that is possible under the current language input. The syntactic structure that receives the most support from the various sources of information will eventually be chosen by means of a competition process (MacDonald, 1994; MacDonald et al., 1994; McRae, Spivey-Knowlton, & Tanenhaus, 1998; Spivey-Knowlton & Sedivy, 1995; Tanenhaus & Trueswell, 1995; Trueswell & Tanenhaus, 1994). To differentiate between these standard models based on *competition* and our own approach based on *optimization* we will refer to the earlier models

as the 'standard' constraint-based models. In the next section we flesh out our model and illustrate how it works.

2. Proposed OT Model: General Mechanism

In OT, inputs are mapped onto outputs by first generating the possible candidates for each input, which is accomplished by the function GEN (short for *generator*), and then selecting the optimal candidate from among them, through the function EVAL (short for *evaluator*). This selection process (or evaluation) consists of the simultaneous application of a hierarchically ordered set of constraints to the candidate outputs. The constraints differ in strength, and crucially, the strongest constraint has absolute dominance over all the weaker, i.e., lower ranked constraints. All constraints are violable, and the candidate that satisfies the total set of constraints best is the *optimal* candidate. The function GEN has two important features: 1) the number of output candidates that it generates is in principle infinite; 2) the output candidates can have any conceivable structure (i.e., syntactic, phonological, semantic, and can in principle even be non-linguistic (e.g., Beaver, 2004)). These features taken together form an essential property of GEN which is called 'freedom of analysis' (Kager, 1999, p. 20).

In our present model, linguistic knowledge *only* resides in the complete set of constraints CON; the GEN component has no linguistic knowledge whatsoever. This precludes the situation where the grammar (= total of linguistic knowledge) would be responsible for generating or restricting the set of output candidates and at the same time for evaluating these candidates. In other words, if the set of candidate outputs were finite, or would have a predefined structure, the grammar would have to apply twice: once during generation of candidates, and once during their evaluation. This would make the model rather

unparsimonious and hard to test, and therefore it is stipulated that the candidate set is infinite in size and consists of all conceivable outputs. Importantly, the number of output candidates is only *virtually* infinite: at the level of the neural processing that underlies OT models, these candidates are not actually physically represented, but can be said to be present as *possible* outcomes (cf. Smolensky & Legendre, 2006). We will discuss these theoretical requirements, and how they may be implemented, in more detail in Section 5.

Our model of sentence processing builds upon findings in two distinct but complementary approaches in OT. In *OT syntax*, the input is a representation of the meaning that a speaker wants to convey to a hearer (cf. Grimshaw, 1997). The function GEN then outputs an unordered and infinite list of grammatical as well as ungrammatical syntactic structures as possible realizations of the speaker's meaning. These candidate structures compete with one another, and depending on the syntactic constraints that apply, one of the candidates will become the optimal, and hence well-formed, candidate. Thus, OT syntax aims to describe how a speaker maps meaning representations onto syntactic structures. OT *semantics*, on the other hand, takes the perspective of the hearer. Here, the input is an overt sample of acoustic material that a hearer receives from a speaker. In this case, GEN generates an unordered and infinite list of possible interpretations of this material (hence its also being called "INT", short for *interpret* (Stevenson & Smolensky, 2006)). Candidate interpretations compete with one another and the optimal candidate is the interpretation for the input form. All other candidates are impossible or unpreferred interpretations for that form. The input form in OT semantics can be a complete sentence but can also be a sentence fragment. If the input is a sentence fragment, the optimal output corresponds to the interpretation that is preferred by the listener at that specific point in the sentence.

In the last few years it has become increasingly clear that speaking and listening are highly interdependent processes (e.g., Blutner, 2000; Blutner, de Hoop & Hendriks, 2006;

Boersma, 1998; Bouma, 2008; Clark, 1996; Hendriks, de Hoop, Krämer, de Swart & Zwarts, 2010). Listeners may take into account the (syntactic) possibilities that speakers have to their disposal to realize their messages, and speakers are sometimes influenced by their knowledge of the listener. Thus, message formulation and comprehension seem to be closely intertwined. In Sections 3.2 and 3.3 we will give examples of constraints that apply to message formulation and comprehension alike. This interrelation of speaker's demands on the one hand and hearer's demands on the other is sometimes also modeled by *bidirectional* optimization (Blutner, 2000; Blutner et al., 2006). For our present purposes we will stick to the unidirectional approach.

So how does our OT model work? Look, for example, at the following table, which in OT-terminology is called a 'tableau':

<< Tableau 1 about here >>

In this tableau constraints are listed from left to right in order of descending strength, so constraint 1 is stronger than constraint 2 (and stronger also than any other constraint that is placed on the right of it), which can be expressed as "Constraint 1 >> Constraint 2 >> ...". Importantly, violation of a stronger constraint is more serious than violation of a weaker constraint. The input is given in the top left-hand corner of the tableau, and candidate outputs are listed in the first column below the input. For expository reasons, only a relevant subset of candidate outputs is present, in the form of possible parses of the input form (e.g., S-coordination, VP-coordination, NP-coordination). One should keep in mind, though, that the actual output is not merely a syntactic structure, but also a semantic - and a pragmatic - interpretation of this structure. That is, terms like S-coordination, VP-coordination and NP-coordination are used here to refer to aspects of interpretation, such as whether an NP must

be interpreted as a subject or an object or whether the subject of the verb preceding the conjunction also is the subject of the verb following the conjunction, rather than to purely structural properties of an expression. The candidates do have a syntactic structure, though, and can be accepted or rejected on the basis of syntactic constraints. To get a flavor of what these candidates might look like, let us look at a very simple input utterance: "dog bites man". The candidate interpretations, each of which consists of a pairing between the given input form and the possible interpretation (<output-interpretation | input-form>), might then be formalized like this:

<Subject=dog; Verb=bites; Object=man | "dog bites man">

<Subject=man; Verb=bites; Object=dog| "dog bites man">

<Subject=bites; Verb=man; Object=dog | "dog bites man">

<Subject=woman; Verb=bites; Object=man | "dog bites man">

etc.

This is of course a very simplified example, and candidate interpretations also carry information regarding event structure, anaphoric relations etc. Basically, we adhere to the 'immediacy of comprehension' hypothesis formulated by Just and Carpenter (1980), and assume that even incomplete sentences are interpreted as completely as possible, up to and including the highest level of discourse.

The success of an enterprise such as ours crucially depends on showing the applicability of the model for every structure in a given language, especially where ambiguity or complexity are concerned. To this end, we will first make an in-depth excursion into one area, coordinated structures, essentially to identify the relevant constraints and establish their hierarchical order on the basis of empirical evidence regarding on-line processing, but also to provide an existence proof that our proposed incremental OT model of interpretation can work. Some confidence that this is indeed the case comes from the pioneering work of

Stevenson and Smolensky (2006) who showed that several instances of syntactic ambiguity resolution can - in principle - be described in terms of an incremental OT model of syntactic competence and performance (cf., Fanselow, Schlesewsky, Cavar, & Kliegl, 1999; Singh, 2002; but see Gibson & Broihier, 1998). Our approach is in many ways very similar to the one described earlier by Stevenson and Smolensky (2006), but crucially extends their model by using constraints from OT *semantics* as well as from OT *syntax*, in order to explain comprehension preferences (cf. de Hoop and Lamers, 2006; see also Lamers & de Hoop, 2005). In other words, the model we propose in this paper takes Stevenson and Smolensky (2006) as its point of departure, but is also different as it is not specifically concerned with *syntactic structure* building based on syntactic constraints, but, crucially with the semantic and pragmatic *interpretation* of an input string on the basis of both syntactic and non-syntactic constraints. As we mentioned before, OT semantics takes the perspective of the hearer, and hence seems very suitable to deal with comprehension preferences. The main reason for including non-structural constraints is that there is no *a priori* reason to assume that initial parsing decisions are exclusively informed by structural factors. Indeed, in the following we will provide empirical evidence suggesting that some garden-path effects are not syntactically motivated at all, but caused, for instance, by pragmatic factors.

We want to make it clear that we aim to use constraints that have already been proposed in the theoretical and empirical literature, and that have received some form of independent support. All other constraints will be viewed as tentative. In addition, it is essential that constraints are formulated as generally as possible, as we want to identify the general principles that are involved in comprehension. Finally, using the same constraints in the same order for every sentence in the language, as required by the OT definition of grammar, will permit us to generate clear predictions. First, in the next section, we will show

that OT is indeed a feasible model of language processing by looking at the processing of (temporarily) ambiguous coordinated structures.

3. Proposed OT Model: Introducing the Constraints Relevant for Coordination

On the basis of a set of studies (e.g., Frazier, 1987a; Frazier & Clifton, 1997; Hagoort, Brown, Vonk, & Hoeks, 2006; Hoeks, 1999; Hoeks, Vonk, & Schriefers, 2002; Hoeks, Hendriks, Vonk, Brown, & Hagoort, 2006; Kaan & Swaab, 2003) regarding the on-line processing of coordination, we will argue that the interpretation of coordinate structures is dependent on a number of constraints of different kinds: 1) pragmatic , 2) discourse-semantic, 3) syntactic, and 4) lexical-semantic. Adopting the framework of OT allows us to formalize this *cross-modular* constraint interaction. Much of the research regarding on-line language processing has focused on the processing of temporarily ambiguous sentences, as readers' preferences at choice points can provide valuable insights into the mechanism of language comprehension. In the research on coordinate structures the emphasis has been on the NP- versus S-coordination ambiguity. For instance, Frazier (1987a) showed in a self-paced reading study that readers prefer NP-coordination over S-coordination in sentences such as (1a) and (1b); slashes indicate how the sentences were divided into segments.

      1a. Pete kissed Marie and / her sister / <u>today</u>. (NP-coordination)

      1b. Pete kissed Marie and / her sister / <u>laughed</u>. (S-coordination)

It was hypothesized that readers prefer to take the ambiguous NP *her sister* as part of the direct object of *kissed* as in (1a). Consequently, they will then run into trouble when reading the final segment of (1b), where the finite verb *laughed* indicates the ambiguous NP is

actually the subject of a conjoined sentence. And indeed, Frazier found significantly longer reading times for *laughed* in (1b) than for *today* in (1a). This finding was replicated in an eye-tracking experiment by Hoeks et al. (2006), who corrected for a number of confounds in Frazier's earlier study. S-coordinated sentences such as (2a) and (2b) were used, the first of which was temporarily ambiguous, whereas the latter served as a control sentence, made unambiguous by inserting a comma after the first object NP (see also Hoeks et al., 2002). Underlined is the critical verb opened which forces an S-coordination reading, where *the photographer* in (2a) is the subject of a conjoined sentence, going against the preferred reading where *the photographer* is part of the direct object of *embraced* (i.e., conjoined with *the designer*).

2a. The model embraced the designer and the photographer opened a bottle of expensive champagne.

2b. The model embraced the designer, and the photographer opened a bottle of expensive champagne.

With these materials it is possible to compare sentences that are identical in sentence meaning, and to compare regions that are identical in length, frequency, and syntactic category. Hoeks et al. (2006) found modest, but reliable evidence for the NP-coordination preference. The same materials were used by Hagoort et al. (2006) in an ERP experiment examining how the brain responds to reading temporarily ambiguous S-coordinations. They found that these sentences evoke a P600-effect, or SPS (i.e., Syntactic Positive Shift) relative to unambiguous control sentences. A P600/SPS is an ERP component generally elicited by ungrammatical sentences (Hagoort, Brown, & Groothusen, 1993; Osterhout & Holcomb, 1992), sentences with an unpreferred syntactic structure (Osterhout & Holcomb, 1992) or

syntactically complex sentences (Kaan, Harris, Gibson, & Holcomb, 2000), although the

P600 has also been reported to occur following semantic violations (e.g., Hoeks, Stowe, &

Doedens, 2004). The size of the P600-effect in the Hagoort et al. experiment was relatively

small, indicating that the processing difficulty as a result of the NP-coordination preference is

rather modest. This same conclusion can be reached from the results of a study by Kaan and

Swaab (2003). Though they did not report the statistical reliability for the comparison that is

of interest here, their figures show a (modest) P600/SPS for the non-preferred S-coordination

(e.g., *The man is painting the house and the garage is already finished*) as compared to an

unambiguous S-coordination with the connective *but* (e.g., *The man is painting the house but

the garage is already finished*).

In sum, all of the studies on the NP- versus S-coordination ambiguity show that when

the ambiguous NP is encountered, readers prefer to interpret it as an argument of the first

main verb (i.e., NP-coordination) instead as the subject of a new clause (i.e., S-coordination).

As a result, temporarily ambiguous S-coordinations give rise to (modest) processing

difficulty. It has generally been ignored, however, that there is an earlier point at which the

sentence is ambiguous, and that is at the connective itself. Before the ambiguous NP is read,

the sentence can continue as an NP- or an S-coordination, but also as a *VP-coordination*, as

shown in sentence (3).

3. The model embraced the designer and laughed.

And indeed, recent evidence from sentence completion studies has shown that language users

strongly prefer to continue a fragment such as (4) as a VP-coordination.

4. The model embraced the designer and …

In about 86 % of all cases coordinated VPs were produced, as opposed to 9 % NP-

coordinations and 5 % S-coordinations (Hoeks et al., 2002, Exp. 1). [1] This outcome suggests

that language comprehenders expect the connective to be followed by a VP, not by an NP.

Only when the NP is actually presented, and VP-coordination is no longer possible, NP-

coordination becomes the preferred structure. This finding provides us with important clues

as to which constraints may be necessary to describe the processing of coordinate structures.

Of course, it is necessary to supplement the observations from the off-line completion study

with solid *on-line* evidence. But the very strong tendency that was found does seem to

indicate a clear preference on the part of the language user.

3.1 OT Constraint (1): A Pragmatic Constraint on Coordination

Why should there be a VP-coordination preference at the connective? According to Hoeks et

al. (2002) the preference for VP-coordination derives from the fact that language users, and

especially readers, must construct their own default 'topic-structure' in the absence of

prosodic or other topic-marking cues. Topic-structure can be loosely defined as describing

the relation between the *topic* of a sentence: the element referring to an entity about which

information is given, and the *information* that is expressed by the sentence (see, e.g.,

Lambrecht, 1994). In VP-coordinations there is only one topic, which is presumed to be the

default and most frequently occurring situation, whereas for instance S-coordinations contain

the additional topic *the photographer* (e.g., (2a)). Having more than one topic, Hoeks et al.

argue, is unexpected and leads to processing difficulty as readers will have to accommodate

an entity that has not been introduced as a second topic in their model of the discourse (e.g.,

Crain & Steedman, 1985; Lambrecht, 1994). Hoeks et al. (2002) provide strong evidence for

this claim by showing in two on-line studies that the processing difficulty is completely

eliminated in a context that makes *both* entities (e.g., both *the model* and *the photographer*)

likely topics of a subsequent sentence (see also Hoeks, Redeker, & Hendriks, 2009). In

addition, studies in the framework of Centering Theory (e.g., Beaver, 2004; Grosz & Sidner,

1986; Grosz, Joshi, & Weinstein, 1995) strongly emphasize the importance of having a single

topic ('backward looking center' in their terms) that provides the link between the current and

the previous sentence. This preference for a single topic was captured in an OT constraint by

Beaver (2004):

UNIQUE TOPIC: With respect to any sentence, there is exactly one discourse entity

which is the topic of that sentence.

Importantly, the UNIQUE TOPIC constraint does not differentiate between VP-coordination and

NP-coordination, as both constructions have but one topic, namely the subject of the sentence

(see Tableau 2).

<< Tableau 2 about here >>

However, language users do seem to prefer VP-coordination over NP-coordination at the

connective. Thus, another constraint must be involved in creating this preference. [2]

3.2 OT Constraint (2): A Discourse-Semantic Constraint on Coordination

We would like to propose that the constraint prohibiting NP-coordination is the discourse-

semantic constraint DO NOT MODIFY, which is a variant of a constraint proposed earlier by

Singh (2002) as "Do not excessively modify any thing or event" (Singh, 2002, p.35) and

belongs to the family of economy constraints in OT (e.g., Legendre et al., 2001).

DO NOT MODIFY: Do not elaborate a thing or an event

If additional information has to be incorporated into the hearer's model of the discourse, this

constraint clearly favors the introduction of a new event to elaboration of a previously

introduced event. In the sentence at hand, the listener therefore prefers VP-coordination as in

(3), where there are two distinct events (e.g., embracing and laughing), to NP-coordination

where there is only one event (e.g., embracing) which is modified by adding another

participant (e.g., the photographer).

Under the analysis we propose, this aversion to modification of events arises because

at the point where the first object NP (e.g., the designer) has been processed, the 'embracing'

event seems sufficiently described from the point of view of the reader, and all thematic roles

of the verb are satisfied by the two available arguments. This stable interpretation would then

be disturbed by the addition of an element that is new, that has not been introduced, and that

somehow has to receive an extra argument role. The effect of DO NOT MODIFY, then, is the

promotion of events that are only minimally elaborated. This constraint is not specific for

coordination: it applies to all kinds of modification and as such it is closely related to the

'principle of referential success', proposed by Crain and Steedman (1985). There, it is

assumed that readers or hearers do not expect modifiers of things or events unless those are

expressly required for unique identification. Many language comprehension studies have

provided support for this principle (Altmann & Steedman, 1988; Ni, Crain, & Shankweiler,

1996; Van Berkum, Brown, & Hagoort, 1999; but see also Clifton & Ferreira, 1989; Mitchell,

Corley, & Garnham, 1992).

To summarize, with the two constraints defined above, we can now describe one step in the incremental comprehension of coordinated structures, and explain how the VP-coordination preference arises at the connective in structures such as (4). Tableau 3 displays how the optimal VP-coordination interpretation is chosen from among the alternatives (we only show the most prominent ones). In OT it is assumed that constraints are hierarchically ordered, that is, from strongest constraint to weakest constraint. However, in this case both orderings, namely UNIQUE TOPIC >> DO NOT MODIFY and DO NOT MODIFY >> UNIQUE TOPIC produce the same optimal candidate, VP-coordination. In such instances, where there is no direct conflict between constraints, more evidence is needed to specify the correct ranking. This is signified by a dashed instead of a solid boundary between the constraints in Tableau 3 (see Anttila & Cho, 1998, for a discussion of 'partial ranking' within OT).

<< Tableau 3 about here >>

3.3 OT Constraint (3): A Syntactic Constraint on Coordination

Now let us consider the situation in which the conjunction *and* is followed by an NP, as in sentence (5).

5. The model embraced the designer and the photographer ...

We have seen that at the time the conjunction *and* is read, VP-coordination is the optimal parse candidate. However, if a VP-coordination parse is also adopted in (5), this entails that the VP must have non-canonical word order. In English and Dutch, VPs do not normally begin with an NP in a non-embedded clause; the canonical word order in main clauses in

these languages is SVO. Because, in (5), the first constituent following the conjunction is an NP, if the coordinate structure is preferably interpreted as VP-coordination the VP must have the non-canonical word order OV rather than VO. This option is ruled out by the same constraints on syntactic structure that also rule out OV word order in the first conjunct in (5). In this specific case it could be the hearer-oriented variant of the constraint STAY (or: "Do not move"; see Grimshaw, 1997; see also Ackema & Neeleman, 1998), which prohibits movement of lexical items:

STAY: Respect canonical word order.

So, the VP-coordination parse violates the constraint STAY if the second conjunct starts with an NP. If the constraints STAY and UNIQUE TOPIC both outrank DO NOT MODIFY, as in Tableau 4, this will account for the empirical observations discussed above that support the NP-coordination preference for (5). This fact also settles the indeterminacy of the ordering of UNIQUE TOPIC and DO NOT MODIFY. If namely the ordering STAY >> DO NOT MODIFY >> UNIQUE TOPIC is assumed, this will yield S-coordination as the optimal parse, which goes against our empirical observations. [3]

<< Tableau 4 about here >>

Note that the optimal parse of (5), corresponding to the NP-coordination interpretation, violates the DO NOT MODIFY constraint, which made NP-coordination sub-optimal in sentence fragment (4) (see Tableau 3). Nevertheless, NP-coordination is optimal in (5) because the competing analyses violate stronger constraints.

Finally, when in a sentence such as (5) the ambiguous NP is followed by a finite verb, as in (6), all options but the S-coordination are rejected by STAY and other, not further specified syntactic constraints, as no NP-coordinated or VP-coordinated sentence can be construed from the current ordered set of words (for convenience we will use STAY as a label for all of those).

6. The model embraced the designer and the photographer laughed ..

<< Tableau 5 about here >>

During the processing of the S-coordinated sentence (6), there are two occasions where there is a shift from one interpretation to another: 1) when the ambiguous NP is read, the preference for VP-coordination shifts to a preference for NP-coordination, as VP-coordination becomes structurally impossible, and NP-coordination does not violate the UNIQUE TOPIC constraint, and 2) on the arrival of the disambiguating verb the NP-coordination reading becomes impossible and the S-coordinated alternative that has long been suboptimal, then becomes the optimal interpretation of the sentence.

Based on the 'Linking Hypothesis' (i.e., linking linguistic competence and performance) proposed by Stevenson and Smolensky (2006), our assumption is that each of these shifts from one interpretation to another gives rise to processing difficulty. We also concur with them in that *any* change brought about by adding linguistic information can be said to cause processing effort or reanalysis: ".. every incremental step in parsing is a process of revising the prior interpretation .." (Stevenson & Smolensky, 2006, p. 854). Thus, there is no qualitative difference, only a quantitative one, between 'just' adding a word, and going from the preferred interpretation at one word to another interpretation at the following word.

Correctly predicting the *amount* of processing difficulty caused by standard linguistic

operations or by different kinds of garden-path structures, however, has proved to be quite

difficult for existing models of sentence processing, and ours is no exception. As to this

problem, Stevenson and Smolensky (2006) make the suggestion that comparing both the

amount and the nature of constraint-violations between the optimal candidate at a given word

and the optimal candidate at the next word will give an estimate of processing effort (see also

Lamers & de Hoop, 2005, who predict qualitatively and quantitatively different ERP effects

from constraint-violation patterns that arise when the processor goes from one optimal

structure to the next). More research is necessary to evaluate this interesting possibility.

In the beginning of this section, we summarized a number of studies showing that

there is indeed processing difficulty at the disambiguating verb of the S-coordinated

sentences that are in focus here. Unfortunately, there is no empirical work explicitly testing

whether there is processing difficulty due to the VP-coordination preference at the ambiguous

NP in structures such as (5). Nevertheless, the OT model that is formulated in Tableau 5

seems to adequately capture all relevant aspects of processing coordinated sentences such as

(6) and predicts that there will be processing difficulty at the ambiguous NP due to the VP-

coordination preference. Future research must determine whether this prediction is borne out.

The constraints on language comprehension that were introduced in this section and

the previous one, DO NOT MODIFY and STAY are strongly tied to language *production* as well.

The constraint STAY, for instance, will prohibit overt or covert movement in production (see,

e.g., Grimshaw, 1997), thus reducing processing costs for the speaker. In comprehension, as

we have seen above, it allows the hearer to identify grammatical relationships by means of

information from word order. The other constraint, DO NOT MODIFY, prohibits modification of

things or events that the speaker has already mentioned. Like STAY, it is an economy

constraint reducing processing costs for the speaker. In comprehension, it has the effect that the hearer closes off the proposition as soon as possible.

Crucially, the two constraints are *violable*, and are restricted in their effects on production by the presence of conflicting constraints as well as by demands from the hearer's perspective. For example, in production STAY can only be satisfied if the resulting structure does not violate a stronger constraint, such as, for instance, the requirement in English and many other languages that questions must be marked as such by a Wh-expression in initial position, which is non-canonical (cf. Ackema & Neeleman, 1998). As a result of this stronger constraint, which helps a hearer to identify questions, a Wh-expression must move, although that means violating STAY. Similarly, the constraint DO NOT MODIFY cannot apply unboundedly in production either. Otherwise, no modifiers - or NP-coordinations - would ever occur! This constraint is restricted in its application by the speaker taking into account the hearer's aim to retrieve the intended meaning. If the speaker would leave out all modifiers in production, the hearer would not be able to recover the meaning they express in comprehension.

3.4 OT Constraint (4): A Lexical-Semantic Constraint on Coordination

An important factor in the interpretation of linguistic utterances that we have not dealt with yet is plausibility. Plausibility can be thought of as involving three broad, interrelated categories of conceptual knowledge: 1) lexical semantic, or 'thematic' knowledge (e.g., how well thematic elements fit their thematic roles), 2) knowledge about the discourse that is presently under consideration, and 3) general knowledge about the world. For our purposes, we will only discuss the role the first kind of plausibility, that is, the one regarding thematic information, plays in our model. According to McRae et al. (1998), a thematic role is ".. the semantic role or mode of

participation played by an entity in the activity or event denoted by the verb" (McRae et al.,

1998, p. 284). Thematic fit, then, is event-specific world knowledge, reflecting the degree to

which the semantic features or an entity fit the requirements of the thematic role it is assigned by

the associated verb. The chances for alternative interpretations to be optimal decline if thematic

fit of an argument is poor given the requirements of the thematic role assigner that is associated

with it. The use of thematic role information in parsing has been studied extensively (e.g.,

Clifton, Traxler, Mohamed, Williams, Morris, & Rayner, 2003; Just & Carpenter, 1992; Ferreira

& Clifton, 1986; McElree & Griffith, 1995; McRae, Feretti, & Amyote, 1997; McRae et al.,

1998; Stowe, 1989; Tanenhaus, Carlson, & Trueswell, 1989; Trueswell, Tanenhaus, & Garnsey,

1994; see also Pickering & Traxler, 1998). One of the best known sentences in this context is

undoubtedly (7), taken from Ferreira and Clifton (1986).

7. The evidence examined by the lawyer turned out to be unreliable.

In this sentence, the verb *examined* is used as a past participle introducing a reduced relative

clause, but it could also be a tensed main verb, which is the generally preferred reading. The first

NP *the evidence*, however, is inanimate and thus a poor AGENT of *examined*, which could lead

to some kind of processing difficulty if a main verb reading is preferred. Indeed, Ferreira and

Clifton found increased reading times for *examined* in sentences such as (7), as compared to

unambiguous controls (e.g., The evidence that was examined by the lawyer ..), indicating that

readers were aware of this fact. However, this did not lead readers to abolish the main clause

reading; they showed as much processing difficulty when reading the disambiguating by-phrase

by the lawyer as when the first NP was an animate entity that easily could fulfill the AGENT

role (e.g., *the defendant*). Trueswell et al. (1994) challenged this finding by pointing out that

some of the 'poor' AGENTS used by Ferreira and Clifton were not that poor at all. For instance,

the car in "The car towed .." can very well play the role of INSTRUMENT in a towing event. With improved materials Trueswell et al. showed that little or no processing difficulty remained in sentences headed by inanimate NPs that were really poor AGENTs (but see Clifton et al., 2003). Thematic information may also play an important role in the processing of coordinated sentences. This can be clearly seen in a sentence fragment such as (8).

8. Jasper sanded the board and the carpenter ...

Here, the poor thematic fit between *carpenter* and *sanded* argues against NP-coordination, which would normally have been the preferred structure (cf. Tableau 3). If on the basis of the thematic fit information S-coordination is assumed instead of NP-coordination, then no processing difficulty is expected to ensue when a disambiguating verb comes in. And indeed, the results of the experiments by Hoeks et al. (2006), using sentences such as 9a, as compared to 9b, indicated that information regarding thematic fit was used very rapidly, and processing difficulty was largely eliminated. A small amount of residual processing difficulty was found, however, compatible with the prediction given in Tableau 6 (i.e., VP-coordination preferred at connective, but S-coordination at ambiguous NP).

9a. Jasper sanded the board and the carpenter repaired ...
9b. Jasper sanded the board, and the carpenter repaired ...

Thus a lexical-semantic factor such as thematic fit is of great influence on the processing of coordination. We will call the associated constraint THEMATIC FIT (adapted from de Hoop and Lamers, 2006; Lamers and de Hoop, 2005)

THEMATIC FIT: A thematic element must meet the requirements of the thematic role that is assigned to it.

Tableau 6 shows how the constraints interact at the time the ambiguous NP *the carpenter* is read; Tableau 7 shows the constraint interaction at the time the disambiguating verb is presented. On the basis of the currently available evidence we cannot decide on the ordering of the constraints THEMATIC FIT and STAY; these do not conflict in the structure at hand and the order in which they appear in the tableaux is therefore inessential.

<< Tableaux 6 and 7 about here >>

Because the preferred analysis at the ambiguous NP is the same as at the following verb (i.e., both times S-coordination), there is no shift from one interpretation to another, and hence no processing difficulty is predicted at the disambiguating verb.

4. The Proposed OT Model Compared to Other Models of Performance

OT and existing constraint-based models are actually very close with respect to their theoretical foundation. Both are based on principles of interacting constraints from all linguistic levels and both assume that syntactic information has no special status, as in, for instance, the garden-path model (e.g., Frazier, 1987a). In fact, OT can be seen as a special case of constraint-satisfaction theory, where instead of numerical strength- and weight-parameters, a hierarchy of constraints is assumed that is characterized by strict domination. Thus, OT and the standard models may make the same predictions in all cases where multiple

constraints interact. If two theories have equally broad empirical coverage, the only thing to decide between them are *meta-theoretical* advantages that we will show our model has. In the following, we will discuss these meta-theoretic advantages that make OT a promising framework for language processing.

As we said, our present OT model of sentence processing has very much in common with the standard constraint-based models proposed by Trueswell and Tanenhaus (1995) and MacDonald et al. (1994). These models were the first in which the interaction of multiple constraints during ambiguity resolution was given a theoretical foundation. In the best-known computer-implementation of the constraint satisfaction process, the so called competition-integration model (McRae et al., 1998; Spivey-Knowlton & Sedivy, 1995; Spivey & Tanenhaus, 1998), syntactic alternatives (typically two) are represented as pre-existing localist nodes in a connectionist network. The nodes representing the alternatives are connected to 'source' nodes representing a variety of information sources: semantic (e.g., thematic fit of a given NP-Verb combination, as estimated by off-line ratings), syntactic (e.g., a general bias for a certain syntactic structure, as estimated from a corpus), pragmatic (e.g., a discourse context biasing towards one of the syntactic alternatives, again estimated from off-line ratings), lexically probabilistic (e.g., the frequency with which a given lexical item is used with a specific argument structures, estimated from a corpus or completion study), but also 'practical' factors such as the possibly disambiguating information that can be gleaned from parafoveal preview during reading. Each of these constraints provides some degree of support for one or both of the syntactic alternatives, depending on a) the *strength of the evidence* that is provided for either of the candidates (e.g., the activation of the source node), and b) on the *weight of the connections* between the knowledge source and the alternative interpretations. For instance, a thematic fit constraint may support one reading with 83% of its activation and the other one with 17%, depending on how well the NPs fit the thematic

requirements of the verbs used in a specific set of sentences (as determined in a pretest). The amount of activation of each constraint node, which represents the amount of evidence it provides for either of the candidates, is transformed into an input activation to both candidate nodes by taking into account the weight of the connection between the source nodes and the candidate nodes. In this implementation, the *weight parameter* reflects the relevance of the constraint for the issue at hand, and is independent of *constraint strength*. To illustrate this independence: a constraint may normally be very relevant and thus have a large weight, but in a specific set of materials there may not be a bias in terms of constraint strength towards either of the alternatives. In that case the constraint will have little or no impact on the disambiguation process, despite the high weight value. During the process of ambiguity resolution, the alternative interpretations receive activation, but also send positive feedback to the constraints, changing the strength with which they support the interpretations etc. This whole process (also called 'normalized recurrence') is repeated until one of the candidates reaches a criterion level of activation, or until a critical amount of time (in terms of processing cycles) has passed.

Though the OT model and the standard models are quite close, and will in many circumstances make identical predictions regarding ambiguity resolution (and sentence processing as a whole) there are also important differences that set them apart. First, standard constraint-based models have no explicit, testable model of linguistic competence. For instance, the competition-integration model (e.g., McRae et al. 1998; Spivey-Knowlton & Sedivy, 1995; Spivey & Tanenhaus, 1998) relies on an unspecified module to produce the syntactic alternatives that enter the competition process. The lexicalist model proposed by MacDonald (1994) is more explicit about how syntactic structure is actually produced, but seems to depend on unspecified sources of syntactic knowledge from outside the lexicon to construct syntactic representations. Some researchers have voiced doubt on whether it is

possible to devise a purely lexically based parser, especially because it may not work well for verb-final languages (e.g., Frazier, 1995). In addition, coordination may present an extra problem for this kind of model, as evidenced by the computational model for a lexical grammar developed by Vosse and Kempen (2000), which is unable to handle coordinated structures (Vosse & Kempen, p. 130). In our OT model, however, the model of competence and the model of performance coincide, as constraints of the grammar *are* processing constraints and vice versa. In other words, the proposed OT model describes how linguistic structures are processed and produced *with one and the same set of constraints*. Of course, our current proposal is just the very first step - by no means do we come close to presenting a full grammar of coordination, let alone of a given language - but this programmatic statement sets it apart from the existing models.

A second disadvantage of standard models is the difficulty they pose in making concrete predictions. The mechanism underlying standard constraint-based models such as the competition-integration model (McRae et al., 1998) is such that specific predictions can only be derived by running the computer simulation, because so many continuous numerical parameters are involved, as we will see below. Running such a simulation, however, is very difficult because of problems with a) identification of the factors that are involved; b) assessment of the *strength* of these factors (i.e., are they strong or weak?); and c) assessment of the *weight* the system assigns to these factors (i.e., regardless of the strength of a factor, is it important or relatively unimportant for the issue at hand?). The first problem is common to many theories of processing, but OT at least has the explicit aim of developing a *general* set of constraints that applies to all sentences in a language. This contrasts with standard models where a custom set of constraints seems to be assembled for each phenomenon. In addition, the mere *number* of proposed constraints (and possibly also the number of syntactic alternatives) will affect the outcome of the simulation in the competition-integration model,

regardless of whether one or the other added constraint is important for resolving the ambiguity in question. For example, introducing a similar but not identical constraint can severely disrupt the modeling process. OT, on the other hand, is easily expandable: the number of constraints does not matter and neither does it matter whether correlated constraints are included. In fact, using OT will allow the researcher to test which of the proposed constraints dominates the other ones. As to point (b), OT aims to apply this general set of constraints in an *invariant order*, so with a predefined hierarchy of strength of constraints for each linguistic phenomenon, whereas assessing the strength of constraints in standard models is rather problematic. For instance, it is unclear at which level of grain one has to analyze a corpus in order to establish possible frequency biases for one or the other syntactic structure. This can be illustrated by looking at two corpus investigations regarding coordinated structures that aimed to find evidence for a frequency bias underlying the NP-coordination preference. First, in a hand parsed sample consisting of a thousand occurrences of the connective 'en' (the Dutch equivalent of *and*) randomly chosen from a newspaper corpus, 46% cases were NP-coordinations, 15% VP-coordinations and 10% S-coordinations (Hoeks et al., 2006). Thus, on this level there appears to be a firm NP-coordination bias emerging from the corpus. However, these so-called 'coarse-grained' measures are very different from the more fine-grained measures that take into account grammatical function (which is done in order to gauge the strength of the frequency factor in sentences such as (2a) that formed the basis for many experiments described in Section 3): in 6% of all cases, NP-coordinations function as grammatical objects, and in 9% of all cases, S-coordinations are found with two different grammatical subjects. Thus, on this level of grain, VP-coordination is the most frequent structure (still 15%), followed by S-coordination (9%), which seems slightly more frequent than NP-coordination (6%). This clearly is not a strong basis to argue for the observed NP-coordination preference: S-coordination is even slightly more frequent

than NP-coordination. And taking the analysis to an even finer grain by looking at definiteness and animacy (cf. Desmet, Brysbaert, & De Baecke, 2001; Jurafsky, 2003) makes the picture change again: less than 1% of the NP-coordinations that are grammatical objects are also both definite and animate, and about 1% of the S-coordinations coordinate clauses with grammatical subjects that are both definite and animate. So choosing this level of analysis is not helpful either if one wants to explain the NP-coordination preference. This shows that the specific level at which a corpus is analyzed crucially determines whether the researcher will find a frequency bias that corresponds to the empirically observed processing preference. One could object that the number of observations was rather small in the finest grain analysis. Therefore, we conducted a second corpus study on a much larger corpus (the automatically parsed *clef* newspaper corpus (Peters, 2001), containing approximately 4,200,000 sentences). We extracted 30,000 instances where 'en' was preceded by a definite NP, and from this set we randomly chose 1,000 cases for manual coding. Here, we found that out of the 243 instances that conformed to our predefined structure (i.e., Subject-Verb-Object-*en*-...), 58% were VP-coordinations, 22% were NP-coordinations and 20% S-coordinations. If animacy was also taken into account, the pattern turned out to be different for the 207 inanimate object NPs (VP-coordination: 52%; NP-coordination: 16%; S-coordinations: 16%), as compared to the 36 animate object NPs (VP-coordination: 6%; NP-coordination: 6%; S-coordinations: 3%). To summarize, as there is no principled way of choosing the 'right' level of analysis, it is very difficult to establish the strength of a probabilistic constraint.

Corpus frequencies may not agree with data from sentence completion studies aimed at finding the strength of a constraint. Again, coordination is a good example, as completion studies show a strong VP-coordination preference at the connective which is reflected in some fine-grained frequency counts but not in others, and certainly not in coarse-grained

frequency counts. See Rayner and Clifton (2002) for an overview of the literature pertaining to the sometimes problematic use of frequency counts and completion data in sentence processing research (cf. Gibson & Schütze, 1999; Gibson, Schütze, & Salomon, 1996; Merlo, 1994; Pickering, Traxler, & Crocker, 2000; Roland & Jurafsky, 2002; but see Desmet, Brysbaert, & De Baecke, 2002; Mitchell, Cuetos, Corley, & Brysbaert, 1995; Swets, De Baecke, Drieghe, Brysbaert, & Vonk, 2006). And to the last point, (c), there is not yet a fool-proof recipe to independently establish valid weights for the different constraints in the simulation of the competition-integration model (see McRae et al., 1998; Spivey-Knowlton & Sedivy, 1995, Spivey & Tanenhaus, 1998, for discussion). OT, then, appears to be more transparent and to have the distinct, if meta-theoretical, advantage of permitting clear predictions, that can sometimes even be derived by hand, because it assumes that the same general set of constraints applies to all structures in the same invariant order. It is thus much stricter than the standard constraint-based models which seem, in a number of ways, too flexible (e.g., in terms of unconstrained number of factors and strength and weight parameters) to make concrete predictions (see also Hoeks, 1999; Narayanan & Jurafsky, 2004, for similar criticisms).

A third disadvantage of the standard models, and especially of lexicalist ones, is their inability to model the interpretation of ungrammatical utterances. It is generally accepted that ungrammatical utterances are rather frequent, especially in spoken language (cf. Antoine, Caelen, & Caillaud, 1994). Because an ungrammatical string cannot be *generated*, either in the form of a pre-fabricated syntactic alternative or as a lexically based syntactic structure, it can never receive an interpretation. This is an important shortcoming and goes against substantial empirical evidence showing that for instance agreement errors, as in *the boy eat an ice-cream*, do not block semantic interpretation (e.g., Gunter & Friederici, 1999; Gunter, Stowe, & Mulder, 1997; Osterhout & Nicol, 1999).

A final criticism concerns the mechanism that standard constraint-based models propose for the constraint interaction process. Syntactic alternatives are assumed to compete with each other on the basis of the evidence that is available for each, until the system settles into some kind of stable state. If all constraints favor one of the alternatives, this stable state is reached very quickly, at no or little processing cost. However, if the alternatives are equally strongly supported, it will take much longer before a stable state is reached, resulting in measurable processing difficulty. Until now the predictions made on the basis of this mechanism have not received unequivocal empirical support, especially when it comes to the processing of sentences that are globally ambiguous. Here, considerable processing difficulty is predicted when the evidence for each of the competing candidates is (approximately) equally strong. However, there is ample evidence that processing these global ambiguities is actually *easier* than processing the unambiguous versions of these sentences (Traxler et al., 1998; Van Gompel et al., 2000; Van Gompel et al., 2005). Let us illustrate this point with an example adapted from Van Gompel et al. (2005).

10a. The <u>governor of the province</u> retiring after the troubles is very rich. (final interpretation: high attachment)

10b. The <u>province of the governor</u> retiring after the troubles is very rich. (final interpretation: low attachment)

10c. The <u>bodyguard of the governor</u> retiring after the troubles is very rich. (globally ambiguous)

In sentences (10a-c), the relative clause 'retiring after the troubles' can either be attached to the first NP of the sentence (high attachment) or to the second NP (low attachment). In the first two sentences, competition between the alternative analyses can be rather short because

plausibility information from the critical word *retiring* disambiguates towards high

attachment in (10a) and to low attachment in (10b). However, if the two alternative analyses

are approximately equally likely, as in the globally ambiguous (10c) (see Traxler et al.,

1998), then constraint-based models such as the competition-integration model predict

considerable processing difficulty. However, Van Gompel et al. (2005) found that globally

ambiguous sentences were significantly easier than either kind of disambiguated sentence.

Furthermore, they showed that globally ambiguous sentences are no harder to process than

completely unambiguous sentences. To account for these findings, Van Gompel and

colleagues propose the Race-Based model of ambiguity resolution, where the syntactic

alternative that is constructed first (on the basis of all available information from the

preceding part of the utterance plus the word category information of the current word) is

adopted by the processor (i.e., it wins the race), which then may be confirmed or

disconfirmed by the semantic or other information contained in the currently processed word.

In the case of globally ambiguous sentences, it does not matter which of the alternatives is

actually chosen, but in both (10a) and (10b), where one of the alternatives is rejected on the

basis of plausibility there will be processing difficulty in a number of instances. Thus,

disambiguated sentences are predicted to be more costly in processing terms than globally

ambiguous structures. Our OT model can explain Van Gompel's data, but in a different way.

The ambiguous part of the sentences used by Van Gompel et al. consist of a complex

constituent containing two definite NPs. Recent research (e.g.,. Farkas & de Swart, 2007;

Hendriks, de Hoop, Krämer, de Swart, & Zwarts, 2010; Van Hout, Harrigan, de Villiers,

2010) has shown that there are non-trivial presuppositions carried by a definite NP. When

readers encounter a definite NP, such as "the bodyguard", they immediately check the

semantic rules for definiteness (Heim, 1982): 1) "Is there a referent for this NP to begin

with?", and 2) "Can this referent be uniquely identified?" In other words, definite NPs carry

the presuppositions of existence and uniqueness, and typically refer to previously introduced referents which are familiar to the reader. The phrase 'the bodyguard of the governor' contains two referring expressions that are both definite but that do not permit unique identification of either of the two referents. Thus, the reader will expect to be given extra information for the unique identification of these entities. This means that an attachment or modification ambiguity is created even before the modifying expression has been encountered, just as proposed by Green and Mitchell (2006), and contra Clifton and Staub (2008), Van Gompel & Pickering (2007), and Levy (2008). Tableau 8a gives a (highly simplified) view on how OT constraint-satisfaction works in a globally ambiguous sentence, whereas Tableau 8b presents the processing of a disambiguated sentence. As mentioned above, we assume that there is an expectation for modification of either the first NP or the second one even before the disambiguating word 'retiring' is encountered. In the ambiguous condition this does not matter, in the disambiguated conditions this will lead to processing difficulty whenever the option chosen at the previous word (indicated by the empty circle in the tableau) turns out to be an implausible AGENT of 'retiring'.

<< Tableaux 8a and 8b about here >>

In response to the findings by Van Gompel et al. (2005), Green and Mitchell (2006) reported a series of computer simulations suggesting that the competition-integration model can also account for the ambiguity advantage. However, a detailed investigation into Green and Mitchell's simulations (Hoeks, Fitz, & Brouwer, submitted) has casted doubt on this conclusion and suggested that the solution Green and Mitchell propose is not psychologically plausible.

One final point we want to make is the following. An important feature of standard constraint-based models, namely their ability to handle probabilistic (lexical) information, may at first glance appear to be rather difficult to emulate in an OT model with strict domination. Information encoding for instance how frequently a verb is used with a specific argument structure, or how frequent a given syntactic structure is in a listeners' linguistic environment may seem hard to capture in the general constraints that figure in OT. As we have seen above, it may not always be clear exactly which probability one needs to use, but one way to handle probabilistic information in an OT framework is to use a constraint that can 'look up' the necessary frequency information from the lexicon, as suggested by Singh (2002). In Singh's proposal, alternative argument-structures (or subcategorizations, alternative meanings, word categories, etc.) for an ambiguous input are represented as a list of possible structures, ordered by frequency. The OT constraint PROBABILITY ACCESS then accesses the position of a given structure in the ordered list, and assigns violations to structures that are not on first position. Thus, probabilistic information can be used in OT decision making. Nevertheless, it may be a more sensible strategy to look for the *cause* of frequency biases, rather than just model them as they appear in a corpus or completion study. For example, Stevenson and Merlo (1997) propose that differences in frequency of usage are most likely caused by differences in thematic and syntactic aspects of the lexical items in question. In the same vein, Argaman and Pearlmutter (2002) strongly argue that the cause for frequency differences must be found at the level of lexical semantic primitives. This should make it possible to model the use of information that appears as probabilistic, but is, in fact, syntactic or semantic in nature.

OT thus has a number of important advantages with regard to standard constraint-based models. Some of these advantages also apply when comparing OT to another class of models that is relevant here. We will call these the 'Bayesian' models, after their proposed

mode of combining probabilistic evidence (Jurafsky, 1996; Narayanan & Jurafsky, 2004; see Brouwer, Fitz, & Hoeks, 2010; Crocker & Brants, 2000, for a different implementation of combining probabilistic evidence). In Bayesian models, lexical entries, but also syntactic rules are represented as mental objects which may have different a priori probabilities or 'resting activations'. During sentence comprehension, these lexical and structural entries are retrieved from memory, and crucially, all structures that are compatible with the input are constructed by the processor. In other words, at any given time many different alternative syntactic structures can be active in parallel, which are activated, or ranked, in accordance to their probability. There is, however, a limit to how improbable a syntactic alternative is allowed to be: if structures are very improbable relative to the other structures, they are pruned away. This can have the consequence that, at a given point in the sentence, the pruned syntactic structure turns out to be the correct structure after all. In that case, the reader is gardenpathed, because that structure is no longer available, and reanalysis must follow.

This kind of model does have an explicit syntactic theory, but most of the other criticisms that were leveled at the standard constraint-based models apply to the Bayesian models as well: 1) it is hard to make predictions for Bayesian models, for one thing because of the considerable difficulty of getting the right probability data from a corpus; 2) a Bayesian model can never account for the fact that language users can understand ungrammatical sentences, as these specific syntactic structures do not exist in the grammar that is generally used. It is certainly possible to use robust parsers that can cope with ungrammatical input, but it is not completely clear how the resulting structural representations can be used for further semantic interpretation in a straightforward way; 3) they make the wrong prediction about the processing of globally ambiguous sentences. Bayesian models assume parallel activation of syntactic alternatives. Thus, they predict that (10a-c) would be equally easy to process, as high attachment and low attachment analyses are

equally probable and thus neither is pruned away. At the arrival of the disambiguating information (e.g., 'retiring'), then, one of the alternatives will be discarded, leaving the other one to be chosen without causing any difficulty. As a final remark, it is hard to see how this kind of model can be extended to include factors such as semantic plausibility or discourse context that are very important for the comprehension of language.

A final class of models we will discuss here are the so-called 'syntax-first' models. In these models it is assumed that processing is not a one-stage parallel phenomenon, but falls into two distinct stages of processing: a first stage in which syntactic structure is built, and a second stage in which syntactic and non-syntactic information are used to construct the interpretation of an utterance (e.g., Frazier, 1987b; Frazier & Clifton, 1996; 1997; Frazier & Rayner, 1982). Though these models are called syntax-first, there is often no explicitly described and testable syntax module. In addition, making predictions for syntax-first models may seem straightforward, but predictions may vary if different grammar formalisms are used (see Crocker, 1992, for discussion). But what is more, syntax-first models are at best only *partial* models of sentence processing as they do not describe how information from the non-syntactic realm is used in constructing an *interpretation* of a sentence, as opposed to merely a structural description. In other words, despite the effort that has gone into modeling reanalysis processes, most notably by Fodor and colleagues (e.g., Fodor & Inoue, 1998), syntax-first models do not get around to model the presumed second phase of their two-stage model where syntactic structure is transformed into meaning. This same criticism applies to models such as those proposed by Phillips (1996) and Weinberg (1999). In all, then, our OT model seems to have several advantages over existing models of sentence comprehension. We will go into some matters of OT architecture before proceeding to the predictions of our model, pertaining also to other structures than coordination.

5. Proposed OT Model: Issues of Implementation

In Section 2 we provided a sketch of the OT architecture that underlies our model. In this section we will discuss this architecture and its associated mechanisms in some more detail. One important issue in this context is the stipulated 'freedom of analysis' for GEN, leading to an infinite number of candidates that can have all kinds of structure. Is this psychologically plausible? And what is more, can it be implemented?

5.1 OT and Connectionism

OT is in essence a hybrid cognitive architecture, combining rule-governed symbolic processing with parallel subsymbolic processing. It is rooted in connectionism, or neural network modeling, where computations are performed by a network of artificial neurons modeled after the human brain (Prince & Smolensky, 1993/2004, 1997). A neural network consists of artificial neurons, or units, and multiple connections between these units. The input to a network consists of a fixed pattern of activation. Activation then flows through the network to construct an output pattern of activation. The neural network thus maps a specific input pattern to a specific output pattern. Crucial for this mapping are the concepts of *harmony* and harmony maximization, or *optimization*. The harmony of a pattern of activation is a measure of its degree of conformity to the connections between the units in the network (Smolensky, 1986). Connections can be excitatory and have a positive weight, or inhibitory and have a negative weight. In addition, the greater the weight of a connection, the greater its importance to the outcome. The connections can be thought of as embodying various constraints, which typically conflict. A pattern of activation that maximizes harmony (or, minimizes energy) is one that optimally balances the demands of all the constraints in the network.

These ideas found their way into linguistics when it was realized that the concept of harmony maximization in neural network modeling could be applied to theories of grammar. The result was a theory called Harmonic Grammar (cf. Legendre, Miyata, & Smolensky, 1990a, 1990b), at the heart of which lies the view that a grammar is a set of violable and potentially conflicting constraints that apply to combinations of linguistic elements. A grammatical structure is then one that optimally satisfies the total set of constraints defined by the grammar. In Harmonic Grammar, as in artificial neural networks, constraints are weighted. Through a process of summation, the overall effect of the total set of constraints can be determined. OT can be seen as a kind of 'restricted' Harmonic Grammar, where the weight of constraints is not formalized by numerical strengths anymore, but solely by a strict priority ranking. But the mechanism by which OT and Harmonic Grammar arrive at an outcome is essentially the same. In fact, it is possible to implement OT in an Harmonic Grammar framework by choosing a specific set of numerical weights (Smolensky & Legendre, 2006).

5.2 The Infinite Candidates of OT

A crucial feature of OT is the stipulation that the candidate set generated by the function GEN is infinite in size. To understand what this means, it is important to make a difference between 1) the *psychological* phenomenon that is modeled, in this case language comprehension, 2) the *formal* model (including a formal algorithm), here our OT model, that aims to describe (aspects of) this psychological phenomenon, and 3) the *computer-implemented* model, which results from implementing the formal model on a computer. In principle, the operations and algorithms proposed in either of these three domains are logically separate and do not have to share any characteristic at all. For instance, when

implementing a formal model on a computer, the implementer must make practical decisions on how to simulate certain aspects of the formal model. As a result, there may exist different implementations of one and the same formal model that use very different procedures, some of which may have a flavor of biological / psychological plausibility (connecting it to the first domain). In the same vein, formal models and the psychological processes that are modeled do not need to show any overlap in processes or algorithms, as long as certain key features and outcomes of the phenomenon in question are adequately captured. Thus, as long as the formal model is adequate, it may propose an infinite number of candidates on a formal level, even if there is no evidence (or logical possibility) that this infinite number of candidates are also psychologically represented during sentence processing.

Nevertheless, it is possible that our OT model (and OT in general) *does* have a lot in common with the real workings of the human brain/mind. By incorporating core concepts from neural network models (which are formal models that due to features such as flow of activation, distributed processing, etc. are considered to have some biological plausibility), OT attempts to provide a formal description of how the subsymbolic level of the brain might be integrated with the symbolic level that seems necessary for the description and explanation of cognitive processes. As we mentioned, OT is a hybrid architecture. Optimization over discrete, symbolic representations at the higher level of the model is inherited from the continuous optimization that takes place at the lower level of neural processing. It is properties such as this one that makes the OT model more similar to the actual mechanism of language comprehension.

But what about the notion of the infinite candidate set, which seems to be psychologically and biologically impossible? Doesn't it make OT less similar to the actual cognitive processes? The reason that we tend to say "no" to this question, is that at the level of neural computation of the formal model, the number of output candidates is only *virtually*

infinite. There is only *one* actual candidate that is considered at a given moment, and that is

the optimal interpretation; in no sense are the alternative structures ever fully represented.

Whereas it is possible to construct an infinite set of output patterns on the basis of differently

ranked constraints, the output patterns corresponding to the possible candidates will not be

actually present at any given time, but are merely *possible* outcomes of different instances of

optimization. Consequently, the system does not require an exhaustive search within an

infinite set of possible outputs. On the contrary, under this view, constraints are seen as

weighted constructors of the optimal candidate, rather than as filters on a set of candidate

outputs (cf. Hagstrom, 1993). This formal description may actually be very similar to what is

going on in the brain.

There are a number of computer implementations of OT that are relevant for the

present discussion, most of which can be found in the field of OT phonology, modeling how

the optimal surface form is chosen for a given underlying form (Tesar, 1994, 1995, 1996;

Ellison, 1994; Frank & Satta, 1998). Some of these OT models are implemented in an explicit

generate-and-test fashion, so without a neural network layer, and may suffer from the

'decidability problem', as argued by Kuhn (2002). This problem refers to the impossibility of

deciding whether a given string is part of the language generated by the grammar, and is due

to the (near to) infinite number of candidates that must actually be generated in this kind of

computer implementation. The decidability problem can be solved, for practical purposes, by

introducing extra conditions on how the constraints are evaluated (see Kuhn, 2002, for

details). According to Kuhn, the decidability problem does not arise when optimizing from a

given form to its optimal meaning, as in our present proposal, or in bidirectional optimization

(Kuhn, 2001, Ch. 6).

In a different vein, the problems that 'infinite candidates' pose for computational

models has been circumvented by Misker and Anderson (2003), who implemented OT in

ACT-R, a hybrid connectionist/symbolic production-system-based architecture. They

modeled the GEN function as finding analogies between the current input and previous

inputs, determining the transformation that alters the analog input into its output, and then

applying this transformation to the current input, in order to obtain a new output. The

candidate set is never represented entirely, only the best candidate so far is stored. A

competing candidate is generated and compared with the current best. Under this approach,

only two candidates are active at any given time.

In our current proposal we have described a formal algorithm for constraint evaluation

at the symbolic level, though we have only given a rather shallow description of the

subsymbolic layer. We cannot at present offer a computer implementation of either layer;

instead, we present a manual simulation in which we have singled out the most likely

candidates for interpretation 'by hand', and where the optimal interpretation was chosen by

applying the proposed constraints, also by hand. Although our model is incomplete and

unimplemented, it is clear that it gives us great predictive power regarding garden paths and

parsing preferences. In fact, the option of a 'manual' simulation mode makes it possible for

other researchers to derive predictions for our OT model, without having to run a

computerized simulation. Having no computer implementation of the subsymbolic level

leaves us unable to explain more dynamic and graded phenomena in language processing,

such as, for instance, word priming, sentence priming, or predicting the actual amount of

processing difficulty for each construction. We can only make predictions on the symbolic

level, which abstracts away from the actual sub-symbolic processing going on at each word.

Exactly which properties should be explained at which level is partly a matter of debate. For

instance, it might be preferable to model all temporary and dynamic aspects of language

processing at the subsymbolic level, thus restricting the symbolic level to general and

permanent forms of knowledge, whereas the subsymbolic level can be taken to represent all

knowledge (temporary and permanent) at the same time. More research in this area is definitely needed.

6. Conclusion

In this paper we propose a model of human sentence processing according to which constraints from the grammar need not be augmented with separate processing restrictions, but rather are processing restrictions themselves. As a result, no distinction needs to be postulated between a competence model (grammar) and a performance model (human parser). This effect is achieved by using Optimality Theory as the competence and performance model. We illustrated the workings of this OT model of human sentence processing by investigating the phenomenon of coordination. Psycholinguistic evidence strongly suggests that the online comprehension of coordinate structures is influenced by constraints from many different information sources: pragmatics, discourse semantics, syntax, and lexical semantics. Adopting the framework of OT allows us to formalize this cross-modular constraint interaction. As we have shown, using the OT framework has many advantages: 1) linguistic competence and performance can be described within one model; 2) as a processing model, it accounts for processing phenomena associated with coordinated structures; 3) it also allows for the interpretation of ungrammatical utterances; and finally 4) it provides very clear and testable predictions. Thus, even though the model that is presented here represents only a first step, it could be a step towards a complete theory of language performance and competence.

Acknowledgments

References

Ackema, P., & Neeleman, A. (1998). Optimal questions. *Natural Language and Linguistic Theory, 16,* 443-490.

Altmann, G. T. M., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition, 30,* 191-238.

Antoine J.-Y., Caelen J., & Caillaud B. (1994). Automatic adaptive understanding of spoken language by cooperation of syntactic parsing and semantic priming. *Proceedings of the ICSLP 1994,* 799-803.

Anttila, A., & Cho, Y.-M. Y. (1998). Variation and change in Optimality Theory. *Lingua, 104,* 31-56.

Argaman, V., & Pearlmutter, N. J. (2002). Lexical semantics as a basis for argument structure frequency biases. In P. Merlo & S. Stevenson (Eds*.), The lexical basis of sentence processing. Formal, computational and experimental issues* (pp. 303-324). Amsterdam, The Netherlands: Benjamins.

Barbosa, P., Fox, D., Hagstrom, P., McGinnis, M., Pesetsky, D. (1998) (Eds.). *Is the best good enough? Optimality and competition in syntax.* Cambridge, MA: MIT Press.

Beaver, D. (2004). The optimization of discourse anaphora. *Linguistics & Philosophy, 27,* 3-56.

Blodgett, A., & Boland, J. E. (1998). *Understanding both in coordinate structures: the OSU subject pool both surprised Allison and Julie!* Poster presented at the eleventh CUNY conference on human sentence processing, March 1998, New Brunswick, NJ.

Blutner, R. (2000). Some aspects of optimality in natural language interpretation. *Journal of Semantics, 17,* 189-216.

Blutner, R., de Hoop, H., & Hendriks, P. (2006). *Optimal Communication.* Stanford, CA: CSLI Publications.

Boersma, P. (1998). *Functional phonology. Formalizing the interactions between articulatory and perceptual drives*. PhD dissertation, University of Amsterdam.

Bouma, G. (2008). *Starting a sentence in Dutch: A corpus study of subject- and object-fronting.* PhD dissertation, University of Groningen.

Bresnan, J. (2000). Optimal syntax. In: J. Dekkers, F. van der Leeuw, & J. van de Weijer (Eds.), *Optimality Theory: phonology, syntax, and acquisition* (pp. 334-385). Oxford, UK: Oxford University Press.

Broekhuis, H, & Vogel, R. (2009). *Optimality Theory and Minimalism: Interface Theories.* Postdam, Germany: Potsdam University Press.

Broekhuis, H, & Vogel, R. (2010). Crash-proof syntax and filters. In: Michael T. Putnam (Ed.) *Crash-proof grammars* (pp. 245-268). Amsterdam, The Netherlands: Benjamins.

Brouwer, H., Fitz, H., & Hoeks, J.C.J. (2010). *Modeling the Noun Phrase versus Sentence Coordination Ambiguity in Dutch: Evidence from Surprisal Theory.* In: J.T. Hale (Ed.), Proceedings of the 2010 Workshop on Cognitive Modeling and Computational Linguistics, ACL2010 (pp. 72-80), Uppsala, Sweden, July 15th 2010.

Chomsky, N. (1957). *Syntactic structures.* The Hague, The Netherlands: Mouton.

Clark, H. H. (1996). *Using language.* Cambridge University Press.

Clifton, C., Jr., & Ferreira, F. (1989). Ambiguity in context. *Language and Cognitive Processes, 4,* SI 77-103.

Clifton, C., Jr., & Staub, A. (2008). Parallelism and competition in syntactic ambiguity resolution. *Language and Linguistics Compass, 2,* 234–250.

Clifton, C., Jr., Traxler, M. J., Mohamed, M. T., Williams, R. S., Morris, R. K, & Rayner, K.

(2003). The use of thematic role information in parsing: Syntactic processing

autonomy revisited. *Journal of Memory and Language, 49,* 317-334.

Crain, S., & Steedman, M. (1985). On not being led up the garden path: The use of context by

the psychological syntax processor. In D. Dowty, L. Karttunen, & A. Zwicky (Eds.),

Natural language parsing: Psychological, computational, and theoretical

perspectives (pp. 320-358). Cambridge, MA: Cambridge University Press.

Crocker, M. W. (1992). *A logical model of competence and performance in the human*

sentence processor. Unpublished doctoral dissertation, University of Edinburgh,

Edinburgh, United Kingdom.

Crocker, M. W., & Brants, T. (2000). Wide-coverage probabilistic sentence processing.

Journal of Psycholinguistic Research, 29, 647-669.

De Hoop, H., Hendriks, P., & Blutner, R. (2007). On compositionality and bidirectional

optimization. *Journal of Cognitive Science, 8(2),* 137-151.

De Hoop, H. & de Swart, H. (2000). Temporal adjunct clauses in Optimality Theory. *Rivista*

di Linguistica, 12, 107-127.

De Hoop, H. & Lamers, M. (2006). Incremental distinguishability of subject and object. In L.

Kulikov, A. Malchukov, & P. de Swart (Eds.), *Case, valency, and transitivity* (pp.

269-290). Amsterdam, The Netherlands: Benjamins.

Desmet, T., Brysbaert, M., & De Baecke, C. (2002). The correspondence between sentence

production and corpus frequencies in modifier attachment. *Quarterly Journal of*

Experimental Psychology, Section A: Human Experimental Psychology, 55, 879-896.

Ellison, T. M. (1994). *Phonological derivation in Optimality Theory.* Proceedings of the 15th

Conference on Computational Linguistics, 1007-1013.

Fanselow, G., Schlesewsky, M., Cavar, D., & Kliegl, R. (1999). *Optimal parsing: Syntactic parsing preferences and optimality theory.* Unpublished manuscript, University of Potsdam, Potsdam, Germany. Available: http://roa.rutgers.edu/files/367-1299/roa-367-fanselow-5.pdf

Farkas, D.F., & de Swart, H. (2007). Article choice in plural generics. *Lingua 117,* 1657-1676.

Ferreira, F., & Clifton, C. (1986). The independence of syntactic processing. *Journal of Memory and Language, 25,* 348-368.

Fodor, J. D., & Inoue, A. (1998). Attach Anyway. In J. D. Fodor & F. Ferreira (Eds.), *Reanalysis in sentence processing* (pp. 101-141). Dordrecht: Kluwer.

Frank, R., & Satta, G. (1998). Optimality Theory and the generative complexity of constraint violability. *Computational Linguistics, 24,* 307-315.

Frazier, L. (1987a). Syntactic processing: Evidence from Dutch. *Natural Language and Linguistic Theory, 5,* 519–559.

Frazier, L. (1987b). Sentence processing: A tutorial review. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp. 601–681). Hillsdale, NJ: Erlbaum.

Frazier, L. (1995). Constraint satisfaction as a theory of sentence processing. *Journal of Psycholinguistic Research, 24,* 437-467.

Frazier, L., & Clifton, C., Jr. (1996). *Construal*. Cambridge, MA: MIT Press.

Frazier, L., & Clifton, C., Jr. (1997). Construal: Overview, motivation, and some new evidence. *Journal of Psycholinguistic Research, 26,* 277–295.

Frazier, L., & Rayner, K. (1982). Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology, 14,* 178-210.

Gáspár, M. (1999). Coordination in Optimality Theory. *Nordic Journal of Linguistics, 22,*
157-182.

Gibson, E., & Broihier, K. (1998). Optimality theory and human sentence processing. In: P.
Barbosa, D. Fox, P. Hagstrom, M. McGinnis, & D. Pesetsky (Eds.), *Is the best good
enough? Optimality and competition in syntax* (pp. 157-191). Cambridge, MA: MIT
Press.

Gibson, E., & Schütze, C. T. (1999). Disambiguation preferences in noun phrase conjunction
do not mirror corpus frequency. *Journal of Memory and Language, 40,* 263–279.

Gibson, E., Schütze, C. T., Salomon, A. (1996). The relationship between the frequency and
the processing complexity of linguistic structure. *Journal of Psycholinguistic
Research, 25,* 59–92.

Green, M. J., & Mitchell, D. C. (2006). Absence of real evidence against competition during
syntactic ambiguity resolution. *Journal of Memory and Language, 55(1),* 1–17.

Grimshaw, J. (1997). Projections, heads and optimality. *Linguistic Inquiry, 28,* 373-422.

Grosz, B., & Sidner, C.L. (1986). Attention, intentions, and the structure of discourse.
Computational Linguistics, 12, 175-204.

Grosz, B., Joshi, A., & Weinstein, S. (1995). Centering: a framework for modeling the local
coherence of discourse. *Computational Linguistics, 21,* 203-226.

Gunter, T. C., & Friederici, A. D. (1999). Concerning the automaticity of syntactic
processing. *Psychophysiology, 36,* 126-137.

Gunter, T. C., Stowe, L. A., & Mulder, G. (1997). When syntax meets semantics,
Psychophysiology, 34, 660-676.

Hagoort, P., Brown, C. M., & Groothusen, J. (1993). The Syntactic Positive Shift (SPS) as an
ERP measure of syntactic processing. *Language and Cognitive Processes, 8,* 439–
483.

Hagoort, P., Brown, C. M., Vonk, W., & Hoeks, J. C. J. (2006). *Processing coordinated*
    structures: evidence from ERPs. Unpublished Manuscript. FC Donders Centre,
    University of Nijmegen, Nijmegen, The Netherlands.

Hagstrom, P. (1993). *On the problem of infinity and Gen-Eval in Optimality Theory.*
    Unpublished manuscript, MIT, Cambridge, MA.

Heim, I. (1982). *The semantics of definite and indefinite Noun Phrases.* University of
    Massachusetts. PhD. diss.; published 1989, New York: Garland.

Hendriks, P., & de Hoop, H. (2001). Optimality Theoretic Semantics. *Linguistics and*
    Philosophy, 24, 1-32.

Hendriks, P., de Hoop, H., Krämer, I., de Swart, H., & Zwarts, J. (2010). *Conflicts in*
    Interpretation. Equinox Publishing, London.

Hendriks, P., & Spenader, J. (2004). A bidirectional explanation of the pronoun interpretation
    problem. In P. Schlenker & E. Keenan (Eds.), *Proceedings of the ESSLLI'04*
    Workshop on Semantic Approaches to Binding Theory, Nancy, France.

Hendriks, P., & Spenader, J. (2005/2006). When production precedes comprehension: An
    optimization approach to the acquisition of pronouns. *Language Acquisition 13:4,*
    319-348.

Hoeks, J. C. J. (1999). *The processing of coordination: semantic and pragmatic constraints*
    on ambiguity resolution. Doctoral Dissertation. University of Nijmegen, Nijmegen,
    The Netherlands. Available: http://www.let.rug.nl/~hoeks/publications/
    Hoeks1999_dis.pdf

Hoeks, J.C.J., Fitz, H., & Brouwer, H. (submitted). *Can the Competition-Integration Model*
    Really Handle the Ambiguity Advantage? Evidence from Computational Modeling.
    Unpublished Manuscript, University of Groningen.

Hoeks, J.C.J., Hendriks, P., Vonk, W., Brown, C.M., & Hagoort, P. (2006). Processing the

    NP- versus S-coordination ambiguity: thematic information does not completely

    eliminate processing difficulty. *Quarterly Journal of Experimental Psychology, 59(9),*

    1581–1599.

Hoeks, J.C.J., Redeker, G., & Hendriks, P. (2009). Fill the gap! Combining pragmatic and

    prosodic information to make gapping easy. *Journal of Psycholinguistic Research, 38,*

    221-235.

Hoeks, J. C. J., Stowe, L. A., & Doedens, L. H. (2004). Seeing words in context: the

    interaction of lexical and sentence level information during reading. *Cognitive Brain*

    Research, 19, 59-73.

Hoeks, J. C. J., Vonk, W., & Schriefers, H. (2002). Processing coordinated structures in

    context: the effect of topic-structure on ambiguity resolution. *Journal of Memory and*

    Language, 46, 99-119.

Jurafsky, D. (1996). A probabilistic model of lexical and syntactic access and

    disambiguation. *Cognitive Science, 20,* 137-194.

Jurafsky, D. (2003). Probabilistic modeling in psycholinguistics: linguistic comprehension

    and production. In: R. Bod, J. Hay, & S. Jannedy (Eds.), *Probabilistic linguistics* (pp.

    39-96). Cambridge, MA: MIT Press.

Just, M. A., & Carpenter, P. A. (1980). A theory of reading: From eye fixations to

    comprehension. *Psychological Review, 87,* 329-354.

Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual

    differences in working memory. *Psychological Review, 98,* 122–149.

Kaan, E., Harris, A., Gibson, E., Holcomb, P. J. (2000). The P600 as an index of syntactic

    integration difficulty. *Language and Cognitive Processes,* 15, 159-201.

Kaan, E., & Swaab, T. Y. (2003). Electrophysiological evidence for serial sentence

processing: a comparison between non-preferred and ungrammatical continuations.

Cognitive Brain Research, 17, 621-635.

Kager, R. (1999). *Optimality Theory.* Cambridge: Cambridge University Press.

Kuhn, J. (2001). *Formal and Computational Aspects of Optimality-theoretic Syntax.* IMS,

Universität Stuttgart.

Kuhn, J. (2002). OT Syntax: Decidability of Generation-based Optimization. *Proceedings of

the 40th Annual Meeting of the Association for Computational Linguistics (ACL),* 48-

55.

Lambrecht, K. (1994). *Information structure and sentence form: Topic, focus, and the mental

representation of discourse referents.* Cambridge, MA: University Press.

Lamers, M., & de Hoop, H. (2005). Animacy information in human sentence processing: an

incremental optimization approach. *Lecture Notes in Artificial Intelligence, 4338,*

158-171.

Legendre, G., Grimshaw, J., & Vikner, S. (2001) (Eds.). *Optimality-Theoretic syntax.*

Cambridge, MA: MIT Press.

Legendre, G., Miyata, Y., & Smolensky, P. (1990a). Harmonic grammar - A formal multi-

level theory of linguistic well-formedness: Theoretical Foundations. In *Proceedings of

the Twelfth Annual Cognitive Science Society Conference,* Hillsdale, NJ: Erlbaum.

Legendre, G., Miyata, Y., & Smolensky, P. (1990b). Harmonic grammar - A formal multi-

level theory of linguistic well-formedness: An application. *Proceedings of the Twelfth

Annual Cognitive Science Society Conference,* Hillsdale, NJ: Erlbaum.

Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition, 106,* 1126–1177.

MacDonald, M. C. (1994). Probabilistic constraints and syntactic ambiguity resolution.

Language and Cognitive Processes, 9, 157-201.

MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. (1994). The lexical nature of

syntactic ambiguity resolution. *Psychological Review, 101,* 676-703.

McCarthy, J. & Prince, A. (1993). *Prosodic morphology: Constraint interaction and*

satisfaction. Technical Report #3 Rutgers University Center for Cognitive Science,

New Brunswick, NJ.

McCarthy, J. (2008). *Doing Optimality Theory.* Malden, MA: Wiley-Blackwell.

McElree, B., & Griffith, T. (1995). Syntactic and thematic processing in sentence

comprehension: Evidence for a temporal dissociation. *Journal of Experimental*

Psychology: Learning, Memory, and Cognition, 21, 134-157.

McRae, K., Feretti, T. R., & Amyote, L. (1997). Thematic roles as verb-specific concepts.

Language and Cognitive Processes, 12, 137-176.

McRae, K.**,** Spivey-Knowlton, M. J., & Tanenhaus, M. K. (1998). Modeling the influence of

thematic fit (and other constraints) in on-line sentence comprehension. *Journal of*

Memory and Language, 38, 283-312.

Merlo, P. (1994). A corpus-based analysis of verb continuation frequencies for syntactic

processing. *Journal of Psycholinguistic Research, 23,* 435-457.

Misker, J. M. V. & Anderson, J. R. (2003) Combining optimality theory and a cognitive

architecture. In F. Detje, D. Dorner, & H. Straub (Eds.), *Proceedings of the Fifth*

International Conference on Cognitive Modeling (pp. 171–76). Universitäts-Verlag

Bamberg.

Mitchell, D. C., Corley, M. M. B., & Garnham, A. (1992). Effects of context in human

sentence parsing: Evidence against a discourse-based proposal mechanism. *Journal of*

Experimental Psychology: Learning, Memory, and Cognition, 18, 69-88.

Mitchell, D. C., Cuetos, F., Corley, M. M. B., & Brysbaert, M. (1995). Exposure-based

    models of human parsing: Evidence for the use of coarse-grained (non-lexical)

    statistical records. *Journal of Psycholinguistic Research, 24,* 469-487.

Narayanan, S., & Jurafsky, D. (2004). *A Bayesian model of human sentence processing.*

    Unpublished manuscript, Berkeley University. Available: http://www.

    icsi.berkeley.edu/%7Esnarayan/newcog.pdf

Ni, W., Crain, S., & Shankweiler, D. (1996). Sidestepping garden paths: Assessing the

    contributions of syntax, semantics, and plausibility in resolving ambiguities.

    Language and Cognitive Processes, 11, 283-334.

Osterhout, L., & Holcomb, P. J. (1992). Event-related brain potentials elicited by syntactic

    anomaly. *Journal of Memory and Language, 31,* 785-806.

Osterhout, L., & Nicol, J. (1999). On the distinctiveness, independence, and time course of

    the brain responses to syntactic and semantic anomalies. *Language and Cognitive*

    Processes, 14, 283-317.

Peters, C. (Ed.) (2001). Cross-Language Information Retrieval and Evaluation. *Lecture Notes*

    in Computer Science, 2069.

Phillips, C. (1996). *Order and structure.* Unpublished doctoral dissertation. MIT, Cambridge,

    MA. Available: http://www.ling.umd.edu/colin/research/papers/dissertation.pdf

Pickering, M. J., & Traxler, M. J. (1998). Plausibility and recovery from garden paths: an

    eye-tracking study. *Journal of Experimental Psychology: Learning, Memory, and*

    Cognition, 24, 940-961.

Pickering, M. J., Traxler, M. J., & Crocker, M. W. (2000). Ambiguity resolution in sentence

    processing: evidence against frequency-based accounts. *Journal of Memory and*

    Language, 43, 447-475.

Prince, A., & Smolensky, P. (1997). Optimality Theory: From neural networks to universal grammar. *Science, 275,* 1604-1610.

Prince, A., & Smolensky, P. (2004). *Optimality Theory. Constraint Interaction in Generative Grammar.* Malden, MA: Blackwell. (Also appeared as Technical Report TR-2, Rutgers Center for Cognitive Science, Rutgers University, New Brunswick, NJ, April 1993).

Rayner, K., & Clifton, C., Jr. (2002). Language comprehension. In D. Medin (Vol. Ed.), *Stevens Handbook of experimental psychology: Vol. 2. Memory and cognitive processes* (3rd ed., pp. 261-316). New York: Wiley.

Roland, D., & Jurafsky, D. (2002). Verb sense and verb subcategorization probabilities. In P. Merlo & S. Stevenson (Eds.), *The lexical basis of sentence processing. Formal, computational and experimental issues* (pp. 325-345). Amsterdam, The Netherlands: Benjamins.

Schachter, P. (1977). Constraints on coordination. *Language, 53,* 86-103.

Sells, P. (2001) (Ed.). *Formal and empirical issues in optimality theoretic syntax.* Stanford, CA: CSLI Publications.

Singh, R. (2002). *A study of optimality theory and the human sentence processing mechanism.* ILLC, University of Amsterdam, Amsterdam, The Netherlands.

Smolensky, P. (1986). Information processing in dynamical systems: foundations of harmony theory. In J. L. McClelland & D. E. Rumelhart (Eds.), *Parallel distributed processing. Explorations in the microstructure of cognition. Volume 1: Foundations* (pp. 194-281). Cambridge, MA: MIT Press.

Smolensky, P., & Legendre, G. (2006). *The harmonic mind: From neural computation to Optimality-Theoretic grammar.* Cambridge, MA: MIT Press.

Spivey-Knowlton, M., & Sedivy, J. (1995). Parsing attachment ambiguities with multiple

constraints. *Cognition, 55,* 227-267.

Spivey, M., & Tanenhaus, M. K. (1998). Syntactic ambiguity resolution in discourse:

modeling the effects of referential context and lexical frequency. *Journal of

Experimental Psychology: Learning, Memory, and Cognition, 24,* 1521-1543.

Stevenson, S., & Merlo, P. (1997). Lexical structure and parsing complexity. *Language and

Cognitive Processes, 12,* 349-399.

Stevenson, S., & Smolensky, P. (2006). Optimality in sentence processing. In: P. Smolensky

& G. Legendre (Eds.), *The Harmonic Mind* (Chapter 19). Cambridge, MA: MIT

Press.

Stowe, L. A. (1989). Thematic structures and sentence comprehension. In G. N. Carlson &

M. K. Tanenhaus (Eds.), *Linguistic structure in language processing* (pp. 319–357).

Dordrecht: Kluwer Academic.

Swets, T., De Baecke, C., Drieghe, D., Brysbaert, M., & Vonk, W. (2006). Relative clause

attachment in Dutch: On-line comprehension corresponds to corpus frequencies when

lexical variables are taken into account. *Language and Cognitive Processes, 21(4),*

453-485.

Tanenhaus, M. K., Carlson, G. N., & Trueswell, J. C. (1989). The role of thematic structures

in interpretation and parsing. *Language and Cognitive Processes, 4,* 211-234.

Tanenhaus, M. K., & Trueswell, J. C. (1995). Sentence comprehension. In J. Miller & P.

Eimas (Eds.), *Handbook of Perception and Cognition, Vol. 11: Speech, Language and

Communication* (pp. 217-262). New York, NY: Academic Press.

Tesar, B. (1994). *Parsing in Optimality Theory: A dynamic programming approach.*

Technical Report CU-CS-714-94, University of Colorado at Boulder.

Tesar, B. (1995). *Computational Optimality Theory.* Unpublished dissertation, University of

    Colorado at Boulder. Available: http://ruccs.rutgers.edu/~tesar/papers/tesardis.pdf

Tesar, B. (1996). Computing optimal descriptions for Optimality Theory grammars with

    context-free position structures. *Proceedings of ACL 34,* 101-107.

Traxler, M. J., Pickering, M. J., & Clifton, C. (1998). Adjunct attachment is not a form of

    lexical ambiguity resolution. *Journal of Memory and Language, 39,* 558-592.

Trueswell, J. C., & Tanenhaus, M. K. (1994). Towards a lexicalist framework for constraint

    based syntactic ambiguity resolution. In: C. Clifton, L. Frazier, & K. Rayner (Eds.),

    Perspectives on Sentence Processing (pp. 155-179). Hillsdale, NJ: Erlbaum.

Trueswell, J. C., Tanenhaus, M. K., & Garnsey, S. M. (1994). Semantic influences on

    parsing: Use of thematic role information in syntactic ambiguity resolution. *Journal of

    Memory and Language, 33,* 285-318.

Van Berkum, J. J. A, Brown, C. M., Hagoort, P. (1999). Early referential context effects in

    sentence processing: Evidence from event-related brain potentials. *Journal of Memory

    and Language, 41,* 147-182.

Van Gompel, P.G., & Pickering, M.J. (2007). Syntactic Parsing (pp. 289-307). In M.G.

    Gaskell and G.T.M. Altmann (Eds.), *Oxford Handbook of Psycholinguistics.* Oxford

    University Press.

Van Gompel, R. P. G., Pickering, M. J., Pearson, J., & Liversedge, S. P. (2005). Evidence

    against competition during syntactic ambiguity resolution. *Journal of Memory and

    Language, 52,* 284-307.

Van Gompel, R. P. G., Pickering, M. J., & Traxler, M. J. (2000). Unrestricted race: a new

    model of syntactic ambiguity resolution. In A. Kennedy, R. Radach, D. Heller, & J.

    Pynte (Eds.), *Reading as a perceptual process* (pp. 621-648). Oxford, UK: Elsevier.

Van Hout, A., Harrigan, K., & de Villiers, J. (2010). Asymmetries in the Acquisition of

   Definite and Indefinite NPs. *Lingua, 120,* 1973-1990.

Vosse, T., & Kempen, G. (2000). Syntactic structure assembly in human parsing: a

   computational model based on competitive inhibition and a lexical grammar.

   Cognition, 75, 105-143.

Weinberg, A. (1999). A minimalist theory of human sentence processing. In: Epstein &

   Hornstein (Eds.), *Working minimalism* (pp. 283-316). Cambridge, MA: MIT Press.

Zeevat, H. (2000). The asymmetry of optimality theoretic syntax and semantics, *Journal of

   Semantics, 17,* 243-262.

Tableau 1: Optimization of a sentence fragment

Input: *fragment*	Constraint 1	Constraint 2	etc.
S-coordination			
VP-coordination			
NP-coordination			

Note. Leftmost column (from top to bottom): input fragment followed by candidate

interpretations; other columns exemplify constraints.

*Tableau 2: Optimization of sentence fragment (4) at the occurrence of **and***

Input: The model embraced the designer **and**	UNIQUE TOPIC
S-coordination	*
● VP-coordination	ok
● NP-coordination	ok

Note. ●=optimal candidate for interpretation of the input; * = constraint violation;

ok=constraint satisfaction.

*Tableau 3: Optimization of sentence fragment (4) at the occurrence of **and***

Input: The model embraced the designer **and**	UNIQUE TOPIC	DO NOT MODIFY
S-coordination	*	ok
● VP-coordination	ok	ok
NP-coordination	ok	*

Note. ●=optimal candidate for interpretation of the input; * = constraint violation;

ok=constraint satisfaction; <u>dashed boundary</u>: ranking between constraints has not been

determined.

*Tableau 4: Optimization of sentence fragment (5) at the occurrence of **the photographer***

Input: The model embraced the designer and **the photographer** ...	STAY	UNIQUE TOPIC	DO NOT MODIFY
S-coordination	ok	*	ok
○ VP-coordination	*	ok	ok
● NP-coordination	ok	ok	*

Note. ○=preferred interpretation at previous word; ●=currently preferred interpretation; * = constraint violation; ok=constraint satisfaction.

*Tableau 5: Optimization of sentence fragment (6) at the occurrence of **laughed***

Input: The model embraced the designer and the photographer **laughed**	STAY	UNIQUE TOPIC	DO NOT MODIFY
● S-coordination	ok	*	ok
VP-coordination	*	ok	ok
○ NP-coordination	*	ok	*

Note. ○= preferred interpretation at previous word; ●=currently preferred interpretation; * = constraint violation; ok=constraint satisfaction.

*Tableau 6: Optimization of sentence fragment (9a) at the occurrence of **the carpenter***

Input: Jasper sanded the board and **the carpenter**	THEMATIC FIT	STAY	UNIQUE TOPIC	DO NOT MODIFY
● S-coordination	ok	ok	*	ok
○ VP-coordination	ok	*	ok	ok
NP-coordination	*	ok	ok	*

Note. ○= preferred interpretation at previous word; ●=currently preferred interpretation; * = constraint violation; ok=constraint satisfaction; <u>dashed boundary</u>: ranking between constraints has not been determined.

*Tableau 7: Optimization of sentence fragment (9a) at the occurrence of **repaired***

Input: Jasper sanded the board and the carpenter **repaired**	THEMATIC FIT	STAY	UNIQUE TOPIC	DO NOT MODIFY
● S-coordination	ok	ok	*	ok
VP-coordination	ok	*	ok	ok
NP-coordination	*	*	ok	*

Note. ○= preferred interpretation at previous word; ●=currently preferred interpretation; * = constraint violation; ok=constraint satisfaction; <u>dashed boundary</u>: ranking between constraints has not been determined.

*Tableau 8a: Optimization of the fragment "The bodyguard of the governor retiring ..." at the occurrence of **retiring***

Input: The bodyguard(NP1) of the governor(NP2) RETIRING	THEMFIT	STAY	UNITOP	DONOTMOD
● NP1-attachment	ok	ok	ok	*
● NP2-attachment	ok	ok	ok	*

Note. ○= preferred interpretation at previous word; ●=currently preferred interpretation; * = constraint violation; ok=constraint satisfaction; <u>dashed boundary</u>: ranking between constraints has not been determined

Tableau 8b: Optimization of the fragment "The province of the governor retiring ..." at the

*occurrence of **retiring***

Input: The province(NP1) of the governor(NP2) RETIRING	THEMFIT	STAY	UNITOP	DONOTMOD
○ NP1-attachment	*	ok	ok	*
● NP2-attachment	ok	ok	ok	*

Note. ○= preferred interpretation at previous word; ●=currently preferred interpretation; * = constraint violation; ok=constraint satisfaction; <u>dashed boundary</u>: ranking between constraints has not been determined

Footnotes

1.   There is evidence that using so-called 'bare plurals' instead of definite NPs as grammatical objects may change the pattern of completion (Blodgett and Boland, 1998). Explaining why this might be the case goes beyond the scope of this paper.

2.   In this example and subsequent ones, we will assume that readers only consider coordination of elements of the same category, that is: S and S, or NP and NP, but not S and NP. However, we believe that this is not so much a constraint on interpretation, but rather an effect of readers taking into account production results. Following Gáspár (1999), we assume sentence production to be subject to the constraint FUSION, which forces duplicate elements to be fused. If this constraint is ranked above the constraints responsible for X-bar structure, the X-bar schema is predicted to be violable. However, violations will only be allowed in order to satisfy FUSION, that is, to yield a coordinate structure. The resulting structures will generally conform to the so-called 'law of coordination of likes' (see, e.g., Chomsky, 1957; Schacter, 1977). If we assume that readers take into account which structures are possible structures in production (as is formalized in bidirectional OT, cf. Blutner, 2000), it follows that they will only consider options where the conjuncts are of the same type in comprehension.

3.   At the ambiguous NP, two constraint orderings produce the desired NP-coordination preference, namely, STAY >> UNIQUE TOPIC >> DO NOT MODIFY and UNIQUE TOPIC >> STAY >> DO NOT MODIFY. At the disambiguating verb, there are three constraint orderings that produce the S-coordination as the optimal candidate, namely STAY >> UNIQUE TOPIC >> DO NOT MODIFY, Stay >> DO NOT MODIFY >> UNIQUE TOPIC, and DO NOT MODIFY >> STAY >> UNIQUE TOPIC. Because OT requires an invariant ordering of constraints, STAY >> UNIQUE TOPIC >> DO NOT MODIFY must be the correct order.